

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Methods and Systems For Synchronizing Data Streams**

Inventor(s):  
Glenn Evans

20051027 01508

10047862-01502

1 **TECHNICAL FIELD**

2 This invention relates generally to processing media content and, more  
3 particularly, to systems and methods for synchronizing media streams.

4  
5 **BACKGROUND**

6 Recent advances in computing power and related technology have fostered  
7 the development of a new generation of powerful software applications. Gaming  
8 applications, communications applications, and multimedia applications have  
9 particularly benefited from increased processing power and clocking speeds.  
10 Indeed, once the province of dedicated, specialty workstations, many personal  
11 computing systems now have the capacity to receive, process and render  
12 multimedia objects (e.g., audio and video content). While the ability to display  
13 (receive, process and render) multimedia content has been around for a while, the  
14 ability for a standard computing system to support true multimedia editing  
15 applications is relatively new.

16 In an effort to satisfy this need, Microsoft Corporation introduced an  
17 innovative development system supporting advanced user-defined multimedia  
18 editing functions. An example of this architecture is described in U.S. Patent No.  
19 5,913,038, issued to Griffiths and commonly owned by the assignee of this  
20 document, the disclosure of which is expressly incorporated herein by reference.

21 In the '038 patent, Griffiths introduced an application program interface  
22 which, when exposed to higher-level development applications, enables a user to  
23 graphically construct a multimedia processing project by piecing together a  
24 collection of "filters" exposed by the interface. The interface described therein is  
25 referred to as a filter graph manager. The filter graph manager controls the data

1 structure of the filter graph and the way that data moves through the filter graph.  
2 The filter graph manager provides a set of object model interfaces for  
3 communication between a filter graph and its application. Filters of a filter graph  
4 architecture implement one or more interfaces, each of which contains a  
5 predefined set of functions, called methods. Methods are called by an application  
6 program or other objects in order to communicate with the object exposing the  
7 interface. The application program can also call methods or interfaces exposed by  
8 the filter graph manager object.

9 Filter graphs work with data representing a variety of media (or non-media)  
10 data types, each type characterized by a data stream that is processed by the filter  
11 components comprising the filter graph. A filter positioned closer to the source of  
12 the data is referred to as an upstream filter, while those further down the  
13 processing chain is referred to as a downstream filter. For each data stream that  
14 the filter handles it exposes at least one virtual pin (i.e., distinguished from a  
15 physical pin such as one might find on an integrated circuit). A virtual pin can be  
16 implemented as an object that represents a point of connection for a unidirectional  
17 data stream on a filter. Input pins represent inputs and accept data into the filter,  
18 while output pins represent outputs and provide data to other filters. Each of the  
19 filters includes at least one memory buffer, and communication of the media  
20 stream between filters is often accomplished by a series of "copy" operations from  
21 one filter to another.

22 A filter graph can have a number of different types of filters, examples of  
23 which include source filters, decoder filters, transform filters, and render filters. A  
24 source filter is used to load data from some source, a decoder filter is used to  
25 decode or decompress a compressed data stream, a transform filter processes and

1 passes data, and a render filter renders data to a hardware device or other locations  
2 (e.g., to a file, etc.).

3 Fig. 1 shows an exemplary filter graph 100 for rendering media content.  
4 Filter graph 100 comprises a number of different filters 104-110 and may or may  
5 not comprise a source 102. A typical filter graph for multimedia content can  
6 include, for example, of graph portion that is dedicated to processing video content  
7 and a graph portion that is dedicated to processing audio content. For example, in  
8 Fig. 1 a source 102 provides content that is typically in compressed form. A  
9 source filter 104 receives the content and then provides the content to one or more  
10 decoder filters for decompression. In this example, consider that filters 106-110  
11 process video content, filters 106a-108a process sub-picture content (such as that  
12 used in Digital Video Data (DVD)), and filters 106b-110b process audio content.  
13 Accordingly, the decoder filters decompress the data and provide the data to a  
14 transform filter (e.g. filters 108-108b) that operates on the data in some way. The  
15 transform filters then provide the transformed data to a corresponding render filter  
16 (e.g. 110, 110b) that then renders the data.

17 Typically, an application program or application 112 provides a means by  
18 which a user can interact with the content that is processed by the filter graph.  
19 Responsive to a user interacting with the application, the application can issue  
20 commands to the source filter 104. Examples of commands can include Run,  
21 Stop, Fast Forward, Rewind, Jump and the like. The source filter receives the  
22 commands and then takes steps to ensure that the commands are executed at the  
23 right time. For example, the source filter 104 typically receives data and provides  
24 timestamps onto data samples that define when the data sample is to be rendered  
25 by the render filters. The source filter then hands the timestamped data sample off



- For the data samples with the input timestamps of 11-20, they wish to have the samples rendered at 5 times the normal rate (i.e. fast forwarded at 5x).

As part of the process that takes place, the decoder filters can adjust the timestamps for the relevant samples so that the samples' output timestamps now comport with the desired playback speeds (i.e. play at 1-1 rate and fast forward at 5x). For example, in order to render the data samples that originally had timestamps of 11-20 (10 timestamps in total) at 5 times the playback rate, those samples will need to be rendered as if they had timestamps of 11 and 12 (i.e. 2 timestamps in total).

So, with this in mind, consider again Fig. 2. For input timestamps of 1-10 there is a one-to-one correspondence between input and output timestamps, meaning that the data samples will be rendered at a normal play rate. Input timestamps of 11-20 will, however, be mapped to output timestamps of 11 and 12 because of the 5x fast forward play rate. Thus, when the render filters receive the data samples with the re-mapped timestamps, the data samples will be rendered in accordance with the desired playback speeds.

Now, in reality, the re-mapping of timestamps can lead to synchronization problems in the following way. Consider, for example, that the individual decoder filters can have different computational models. That is, the different decoder filters might be provided from different vendors. Accordingly, the different computational models may perform computations for purposes of re-mapping time stamps differently. Specifically, the computational models may perform rounding operations differently. Because of this, the re-mapped timestamps can vary as

1 between data samples that should for all practical purposes be rendered together.  
2 This can manifest itself in some different ways. For example, the audio that  
3 accompanies the video may lag just enough to be annoying. Additionally, sub-  
4 pictures such as video overlays may be overlaid at the wrong time. Thus, the user  
5 experience can be degraded.

6 Products utilizing the filter graph have been well received in the market as  
7 it has opened the door to multimedia editing using otherwise standard computing  
8 systems. Yet, there continues to be a need to improve filter graph technology and  
9 further enhance the user experience, or at least not degrade it.

10 Accordingly, this invention arose out of concerns associated with providing  
11 improved methods and systems for synchronizing timestamped data streams and,  
12 in particular, timestamped data streams associated with filter graphs.

### 13 14 SUMMARY

15 Methods and systems are provided for synchronizing various time-stamped  
16 data streams. The data streams can be synchronized to another data stream or to a  
17 point of reference such as a reference clock. Synchronization can take place  
18 periodically or in accordance with a defined tolerance which, if equaled or  
19 exceeded, can be used to trigger a synchronization process.

20 In one embodiment, synchronization processing takes place in association  
21 with a filter graph comprising multiple filters. The filter graph is configured to  
22 process multiple timestamped data streams for rendering the data streams in  
23 accordance with data stream timestamps. A synchronization module is provided  
24 and is associated with the filter graph. The synchronization module is configured  
25 to query individual filters of the filter graph to ascertain input timestamp-to-output

1 timestamp mappings. The module then computes adjustments that are to be made  
2 to output timestamps in order to synchronize the data streams, and then instructs  
3 queried filters to adjust their output timestamps in accordance with its adjustment  
4 computations.

## 6 **BRIEF DESCRIPTION OF THE DRAWINGS**

7 Fig. 1 is a diagram of an exemplary conventional filter graph.

8 Fig. 2 is a graph that is useful in understanding various concepts associated  
9 with the described embodiments.

10 Fig. 3 is a block diagram that illustrates an exemplary computer system that  
11 can be used to implement various embodiments described below.

12 Fig. 4 is a diagram of an exemplary filter graph and is useful in  
13 understanding various concepts associated with the described embodiments.

14 Fig. 5 is a graph that describes input/output timestamp mappings.

15 Fig. 6 is a diagram of an exemplary filter graph and synchronization  
16 module in accordance with one embodiment.

17 Fig. 7 is a graph that describes input/output timestamp mappings associated  
18 with the Fig. 6 filter graph.

19 Fig. 8 is a flow diagram that describes steps in a method in accordance with  
20 one embodiment.

21 Fig. 9 is a flow diagram that describes steps in a method in accordance with  
22 one embodiment.

23 Fig. 10 is a graph that describes input/output timestamp mappings  
24 associated with another embodiment.  
25



## DETAILED DESCRIPTION

### Overview

Methods and systems are provided for synchronizing various time-stamped data streams. The data streams can be synchronized to another data stream or to a point of reference such as a reference clock. Synchronization can take place periodically or in accordance with a defined tolerance which, if equaled or exceeded, can be used to trigger a synchronization process.

### Exemplary Computing Environment

Fig. 3 illustrates an example of a suitable computing environment 300 on which the system and related methods for processing media content may be implemented.

It is to be appreciated that computing environment 300 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the media processing system. Neither should the computing environment 300 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 300.

The media processing system is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the media processing system include, but are not limited to, personal computers, server computers, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers,

1 mainframe computers, distributed computing environments that include any of the  
2 above systems or devices, and the like.

3 In certain implementations, the system and related methods for processing  
4 media content may well be described in the general context of computer-  
5 executable instructions, such as program modules, being executed by a computer.  
6 Generally, program modules include routines, programs, objects, components,  
7 data structures, etc. that perform particular tasks or implement particular abstract  
8 data types. The media processing system may also be practiced in distributed  
9 computing environments where tasks are performed by remote processing devices  
10 that are linked through a communications network. In a distributed computing  
11 environment, program modules may be located in both local and remote computer  
12 storage media including memory storage devices.

13 In accordance with the illustrated example embodiment of Fig. 3 computing  
14 system 300 is shown comprising one or more processors or processing units 302, a  
15 system memory 304, and a bus 306 that couples various system components  
16 including the system memory 304 to the processor 302.

17 Bus 306 is intended to represent one or more of any of several types of bus  
18 structures, including a memory bus or memory controller, a peripheral bus, an  
19 accelerated graphics port, and a processor or local bus using any of a variety of  
20 bus architectures. By way of example, and not limitation, such architectures  
21 include Industry Standard Architecture (ISA) bus, Micro Channel Architecture  
22 (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association  
23 (VESA) local bus, and Peripheral Component Interconnects (PCI) bus also known  
24 as Mezzanine bus.

25

1 Computer 300 typically includes a variety of computer readable media.  
2 Such media may be any available media that is locally and/or remotely accessible  
3 by computer 300, and it includes both volatile and non-volatile media, removable  
4 and non-removable media.

5 In Fig. 3, the system memory 304 includes computer readable media in the  
6 form of volatile, such as random access memory (RAM) 310, and/or non-volatile  
7 memory, such as read only memory (ROM) 308. A basic input/output system  
8 (BIOS) 312, containing the basic routines that help to transfer information  
9 between elements within computer 300, such as during start-up, is stored in ROM  
10 308. RAM 310 typically contains data and/or program modules that are  
11 immediately accessible to and/or presently be operated on by processing unit(s)  
12 302.

13 Computer 300 may further include other removable/non-removable,  
14 volatile/non-volatile computer storage media. By way of example only, Fig. 3  
15 illustrates a hard disk drive 328 for reading from and writing to a non-removable,  
16 non-volatile magnetic media (not shown and typically called a "hard drive"), a  
17 magnetic disk drive 330 for reading from and writing to a removable, non-volatile  
18 magnetic disk 332 (e.g., a "floppy disk"), and an optical disk drive 334 for reading  
19 from or writing to a removable, non-volatile optical disk 336 such as a CD-ROM,  
20 DVD-ROM or other optical media. The hard disk drive 328, magnetic disk drive  
21 330, and optical disk drive 334 are each connected to bus 306 by one or more  
22 interfaces 326.

23 The drives and their associated computer-readable media provide  
24 nonvolatile storage of computer readable instructions, data structures, program  
25 modules, and other data for computer 300. Although the exemplary environment

1 described herein employs a hard disk 328, a removable magnetic disk 332 and a  
2 removable optical disk 336, it should be appreciated by those skilled in the art that  
3 other types of computer readable media which can store data that is accessible by a  
4 computer, such as magnetic cassettes, flash memory cards, digital video disks,  
5 random access memories (RAMs), read only memories (ROM), and the like, may  
6 also be used in the exemplary operating environment.

7 A number of program modules may be stored on the hard disk 328,  
8 magnetic disk 332, optical disk 336, ROM 308, or RAM 310, including, by way of  
9 example, and not limitation, an operating system 314, one or more application  
10 programs 316 (e.g., multimedia application program 324), other program modules  
11 318, and program data 320. In accordance with the illustrated example  
12 embodiment of Fig. 3, operating system 314 includes an application program  
13 interface embodied as a render engine 322. As will be developed more fully  
14 below, render engine 322 is exposed to higher-level applications (e.g., 316) to  
15 automatically assemble filter graphs in support of user-defined development  
16 projects, e.g., media processing projects. Unlike conventional media processing  
17 systems, however, render engine 322 utilizes a scalable, dynamically  
18 reconfigurable matrix switch to reduce filter graph complexity, thereby reducing  
19 the computational and memory resources required to complete a development  
20 project. Various aspects of the innovative media processing system represented by  
21 a computer 300 implementing the innovative render engine 222 will be developed  
22 further, below.

23 Continuing with Fig. 3, a user may enter commands and information into  
24 computer 300 through input devices such as keyboard 338 and pointing device 340  
25 (such as a "mouse"). Other input devices may include a audio/video input

1 device(s) 353, a microphone, joystick, game pad, satellite dish, serial port, scanner,  
2 or the like (not shown). These and other input devices are connected to the  
3 processing unit(s) 302 through input interface(s) 342 that is coupled to bus 306,  
4 but may be connected by other interface and bus structures, such as a parallel port,  
5 game port, or a universal serial bus (USB).

6 A monitor 356 or other type of display device is also connected to bus 306  
7 via an interface, such as a video adapter 344. In addition to the monitor, personal  
8 computers typically include other peripheral output devices (not shown), such as  
9 speakers and printers, which may be connected through output peripheral interface  
10 346.

11 Computer 300 may operate in a networked environment using logical  
12 connections to one or more remote computers, such as a remote computer 350.  
13 Remote computer 350 may include many or all of the elements and features  
14 described herein relative to computer 300 including, for example, render engine  
15 322 and one or more development applications 316 utilizing the resources of  
16 render engine 322.

17 As shown in Fig. 3, computing system 300 is communicatively coupled to  
18 remote devices (e.g., remote computer 350) through a local area network (LAN)  
19 351 and a general wide area network (WAN) 352. Such networking environments  
20 are commonplace in offices, enterprise-wide computer networks, intranets, and the  
21 Internet.

22 When used in a LAN networking environment, the computer 300 is  
23 connected to LAN 351 through a suitable network interface or adapter 348. When  
24 used in a WAN networking environment, the computer 300 typically includes a  
25 modem 354 or other means for establishing communications over the WAN 352.

20047862-01502

1 The modem 354, which may be internal or external, may be connected to the  
2 system bus 306 via the user input interface 342, or other appropriate mechanism.

3 In a networked environment, program modules depicted relative to the  
4 personal computer 300, or portions thereof, may be stored in a remote memory  
5 storage device. By way of example, and not limitation, Fig. 3 illustrates remote  
6 application programs 316 as residing on a memory device of remote computer  
7 350. It will be appreciated that the network connections shown and described are  
8 exemplary and other means of establishing a communications link between the  
9 computers may be used.

### 11 Exemplary Embodiment

12 For purposes of understanding various principles upon which the various  
13 inventive embodiments are based, consider Fig. 4.

14 There, a filter graph 400 is shown and is similar to filter graph 100 in Fig.  
15 1. Assume in this example, that each of the decoder filters 106, 106a and 106b is  
16 slightly computationally different in that output timestamps are assigned in a  
17 slightly different way. For example, assume that the application 112 has indicated  
18 to the source filter 104 that the user wishes to have the data streams rendered at 2x  
19 the playback rate. Assume also that because of the computational differences of  
20 the various decoders, timestamps are re-mapped in such a way that the video  
21 stream associated with decoder filter 106 will be rendered at 2.1x the playback  
22 speed (i.e. slightly faster); the sub-picture stream will be rendered at 1.9x the  
23 playback speed (i.e. slightly slower); and the audio stream will be rendered at 2.0x  
24 the playback speed (i.e. the correct speed). As will be appreciated, these streams  
25 will, over time, tend to drift relative to one another.

As an example, consider Fig. 5 which shows a graph that illustrates the mapping of the input timestamps to output timestamps for each of the Fig. 4 decoders. Specifically, line 502 comprises the mapping for the decoder associated with the video stream (i.e. decoder 106 in Fig. 4), line 504 comprises the mapping for the decoder associated with the audio stream (i.e. decoder 106b in Fig. 4), and line 506 comprises the mapping for the decoder associated with the sub-picture stream (i.e. decoder 106a in Fig. 4). These lines or curves should ideally, without any drift, lie on top of each other. That is, without any drift, the input timestamps for each input timestamp value should map to the same output timestamp value. Unfortunately, because of the drift, this does not occur. For example, notice that for an input timestamp of 10, the output timestamp for each of the streams is different.

Consider now Fig. 6. There, a synchronization module 600 is provided. The synchronization module can be implemented in any suitable hardware, software, firmware or combination thereof. In the illustrated example, the synchronization module is implemented in software.

The synchronization module is configured to periodically query individual filters and instruct the filters to adjust the output timestamps of individual data samples so that the data streams are synchronized. In the present example, module 600 queries each of the decoder filters and then instructs the decoder filters to adjust the output timestamps for synchronizing the data streams. In this particular embodiment, the synchronization module comprises a filter query module 602 that queries the individual filters, and a stream adjustment module 604 that computes the adjustments that should be made to the output timestamps.

1        One solution to synchronizing the individual data streams is to ascertain the  
2        current input timestamp and assume that all decoder filters are at the current input  
3        time stamp. The decoder filters can then be queried as to their output timestamp  
4        mappings for the assumed current input timestamp mapping. When the decoder  
5        filters respond with their corresponding output timestamp, the furthest output  
6        timestamp can be ascertained and then the decoder filters that do not correspond to  
7        the decoder filter having the furthest output timestamp can be instructed to start  
8        assigning output timestamps at a value equal to the furthest output timestamp.

9        Consider, for example, Fig. 5. When the decoder filters are queried, they  
10       will each respond with their current output timestamp. However, because the  
11       streams are continuously being processed, the assumption that the current input  
12       timestamp is at 10 is not entirely accurate. For example, when decoder filter 106  
13       (Fig. 6) is queried, the input timestamp may well be 10. Thus, for decoder 106  
14       this is a good assumption. However, because of the serial nature of the querying  
15       and the advancing time, when decoder filter 106a is queried, the corresponding  
16       input timestamp will likely not be 10, but rather might be 10.1. Thus, for decoder  
17       106a, the assumption that the current input timestamp is 10 is not an accurate  
18       assumption. Similarly, when the decoder filter 106b is queried, the input  
19       timestamp may actually be 10.2. Thus, for decoder 106b, the assumption that the  
20       current input timestamp is 10 is not an accurate assumption.

21       Thus, while this approach may bring the data streams into closer  
22       synchronization, this is not the best as its underlying assumption concerning the  
23       current input timestamp is not accurate with respect to all of the decoders.  
24  
25



1 Consider now Fig. 6 in connection with Fig. 5. Because of the real time  
2 nature of the environment in which the querying takes place, the data streams are  
3 simultaneously being processed while the querying takes place.

4 In accordance with one embodiment, each of the decoder filters is queried  
5 to ascertain the current input timestamp and the current output timestamp. In  
6 addition, if the actual playback rate of the decoder is not known, each decoder can  
7 be queried for its playback rate. Once this information is ascertained,  
8 synchronization module 600 can compute an output timestamp for a specific input  
9 timestamp and then instruct one or more of the decoders to synchronize their  
10 output timestamps to the computed output timestamp.

11 For example, module 600 can query the individual decoder filters in a serial  
12 fashion. For example, the module 600 might query decoder filter 106 first, and  
13 then decoder filter 106a and then decoder filter 106b. This is diagrammatically  
14 shown in the graph of Fig. 5. There, notice that at a time that corresponds to input  
15 timestamp 10, decoder filter 106 is queried to provide its current input timestamp  
16 and the corresponding output timestamp that is associated with input timestamp  
17 10. Ideally, we assume the decoder would map an input timestamp of 10 to an  
18 output timestamp of 100. Because decoder filter 106 is slightly faster than the  
19 actual requested playback speed it responds with a value of 95. Likewise, at the  
20 next query time (which is shortly after the first query time and which corresponds  
21 to an input timestamp of 10.1), decoder filter 106a is queried to provide its current  
22 input timestamp and the corresponding output timestamp that is associated with  
23 input timestamp 10.1. Because decoder filter 106a is slightly slower than the  
24 actual requested playback speed, it responds with a value of 110. Likewise, at the  
25 next query time (which is shortly after the second query time and which

corresponds to an input timestamp of 10.2), decoder filter 106b is queried to provide its current input timestamp and the corresponding output timestamp that is associated with input timestamp 10.2. Because decoder filter 106b is synchronized with the actual requested playback speed, it responds with a value of 100.1 (see table for computation). Thus, the table below summarizes the mappings of current input timestamps to output timestamps. Note additionally that if the playback rates of the decoders are not known, the decoders can be queried for their playback rates.

Decoder	Input Timestamp	Output Timestamp	Rate	Output Timestamp At 10.2 Output +(output-10.2)/rate
Decoder 106	10	95	2.1	$95 + 0.2 / 2.1x = 95.0952$
Decoder 106a	10.1	110	1.9	$110 + 0.1 / 1.9x = 110.0526$
Decoder 106b	10.2	100.1	2.0	$100.1 + 0 / 2.0x = 100.1$

In accordance with one embodiment, once the decoder filters have been queried and have responded with their individual mappings, the synchronization module 600 can extrapolate each of the lines characterizing the timestamp mappings to a defined point corresponding to a common input timestamp. Corrections can then be calculated and the decoders can be instructed to synchronize their output timestamp mappings accordingly.

As an example of how this can be done, consider Fig. 7 which shows a mapping of input timestamps to output timestamps generally at 700. First notice that in this example the line that characterizes each of the mappings of input timestamps to output timestamps can be characterized by the classic line equation

1  $y=mx + b$ . Here, the  $y$  variable represents the output timestamp, the  $x$  variable  
2 represents the input timestamp and the slope  $m$  represents the playback rate, and  $b$   
3 is a constant.

4 In this specific example, each of the lines characterizing the mappings of  
5 input timestamp to output timestamp is extrapolated, if necessary, to the largest  
6 value of input timestamp that was returned by the query. In this example, and  
7 from the table above, the largest input timestamp value that was returned as a  
8 result of the query of decoder filters is 10.2. Accordingly, lines 502 and 506 are  
9 extrapolated to the input timestamp of 10.2. Notice that the input timestamp of  
10 10.2 is represented vertically by the dashed line extending upward from the value  
11 of 10.2. Notice also that the extrapolated portion of each of lines 502, 506 is  
12 respectively shown at 502a and 506a.

13 Once the individual lines have been extrapolated, a skip parameter can be  
14 calculated. The skip parameter represents a value that can be used to synchronize  
15 the output timestamps of the various decoders. In this example, the skip value is  
16 computed by taking the difference between the largest output timestamp value for  
17 the given input timestamp value and the output timestamp value for the line  
18 characterizing the decoder mappings for the given decoders for the given input  
19 timestamp. As an example, consider again Fig. 7. There, the skip value for line  
20 502 is computed by taking the difference between 110.0526 (i.e. the largest output  
21 timestamp value for the given input timestamp value of 10.2) and 95.0952 (i.e. the  
22 output timestamp value for lines 502 at the input timestamp value of 10.2) to  
23 provide a skip value of 14.9574. Likewise, the skip value for line 504 is computed  
24 by taking the difference between 110.0526 (i.e. the largest output timestamp value  
25 for the given input timestamp value of 10.2) and 100.1 (i.e. the output timestamp

1 value for lines 504 at the input timestamp value of 10.2) to provide a skip value of  
2 9.9526.

3 Next, individual decoders are instructed to jump their output timestamps by  
4 their individual skip values for the corresponding input timestamp. Here, for an  
5 input timestamp of 10.2, the individual decoders would be instructed to add their  
6 associated skip value to their output timestamp. This has the desirable effect of  
7 adjusting the ends of each of lines 502, 504 upwardly to coincide with the end of  
8 line 506. Hence, the data streams are brought back into a desirable level of  
9 synchrony.

10 The reader should appreciate that the skip value can be calculated relative  
11 to any desirable common input timestamp. In this particular example, the input  
12 time stamp of the last-queried decoder was used. This need not, however, be the  
13 case. For example, the process can select a particular input timestamp in the  
14 future, say 10.5, and extrapolate all of the lines characterizing the mappings to  
15 10.5. Then, all of the decoders can be instructed to jump by the computed skip  
16 value when the input timestamp value corresponds to 10.5.

17 It should also be noted that the above-described process can be repeated  
18 periodically to ensure that the data streams remain synchronized at a desired level  
19 of synchrony. It should also be appreciated that a skip value tolerance can be  
20 defined and the synchronization process can be performed any time that any of the  
21 decoder skip values exceed or equal the skip value tolerance. For example,  
22 assume that a skip value tolerance of 10 is defined. In this case, the mappings of  
23 input timestamps to output timestamps can be monitored for each of the decoders.  
24 This is diagrammatically analogous to monitoring each of the lines 502, 504, and  
25

1 506. Then, any time a skip value for any of the lines equals or exceeds the skip  
2 value tolerance, the synchronization process can be performed.

3 Fig. 8 is a flow diagram that illustrates steps in a method in accordance  
4 with one embodiment. The method can be implemented in any suitable hardware,  
5 software, firmware or combination thereof. In the illustrated example, the method  
6 is implemented in software. The method can be implemented by a  
7 synchronization module such as module 600 of Fig. 6.

8 Step 800 queries one or more filters for their input/output timestamp  
9 mappings. Additionally, if the individual playback rates for the filters are not  
10 known, then step 800 can also query for the playback rates. Any suitable filter can  
11 be queried. In the particular example above, the decoder filters are queried. If, in  
12 some systems, the decoder filters are not the filters that perform the input/output  
13 timestamp mappings, then the filters that perform those mappings can be queried.  
14 Responsive to receiving responses from the queried filters, step 802 extrapolates  
15 lines characterizing the mappings to a selected input timestamp value. The  
16 extrapolation can be accomplished using any suitable extrapolation function. For  
17 example, in the above examples the extrapolation was a linear extrapolation. It is  
18 possible, however, for the lines that characterize the mappings to be non-linear. In  
19 this case, the extrapolation can be non-linear. Additionally, the selected input  
20 timestamp value to which such lines are extrapolated can comprise any desirable  
21 value. For example, the selected input timestamp value can comprise a current  
22 input timestamp value for one of the filters (as in the Fig. 5 example). Alternately,  
23 the current input timestamp value can comprise a future input timestamp value.  
24 Once the lines are extrapolated, step 804 calculates skip values for one or more of  
25 the lines. The skip values represent a value by which the output timestamps for a

1 given filter are to be corrected for the selected input timestamp value. One  
2 example of how skip values can be calculated is given above.

3 Once the skip values are calculated for the individual filters, step 806  
4 provides instructions to synchronize the data streams based on the calculated skip  
5 values. In the above example, this was accomplished by instructing the filters to  
6 skip their output timestamps ahead by an associated skip value, for a selected input  
7 timestamp. Step 806 can then return to step 800 and the process can be  
8 periodically repeated to maintain the data streams in synchrony.

9 Fig. 9 is a flow diagram that illustrates steps in a method in accordance  
10 with one embodiment. The method can be implemented in any suitable hardware,  
11 software, firmware or combination thereof. In the illustrated example, the method  
12 is implemented in software. The method can be implemented by a  
13 synchronization module such as module 600 of Fig. 6.

14 In this process, a skip value tolerance is defined and synchronization  
15 processing is performed whenever the data streams become unsynchronized  
16 enough to meet or exceed the skip value tolerance.

17 Accordingly, step 900 defines a skip value tolerance. Step 902 queries one  
18 or more filters for their input/output timestamp mappings. Additionally, if the  
19 individual playback rates for the filters are not known, then step 902 can also  
20 query for the playback rates. Any suitable filter can be queried. In the particular  
21 example above, the decoder filters are queried. If, in some systems, the decoder  
22 filters are not the filters that perform the input/output timestamp mappings, then  
23 the filters that perform those mappings can be queried. Responsive to receiving  
24 responses from the queried filters, step 904 extrapolates lines characterizing the  
25 mappings to a selected input timestamp value. The extrapolation can be

1 accomplished using any suitable extrapolation function. For example, in the  
2 above examples the extrapolation was a linear extrapolation. It is possible,  
3 however, for the lines that characterize the mappings to be non-linear. In this case,  
4 the extrapolation can be non-linear. Additionally, the selected input timestamp  
5 value to which such lines are extrapolated can comprise any desirable value. For  
6 example, the selected input timestamp value can comprise a current input  
7 timestamp value for one of the filters (as in the Fig. 5 example). Alternately, the  
8 current input timestamp value can comprise a future input timestamp value. Once  
9 the lines are extrapolated, step 906 calculates skip values for one or more of the  
10 lines. The skip values represent a value by which the output timestamps for a  
11 given filter can be corrected for the selected input timestamp value. One example  
12 of how skip values can be calculated is given above.

13 Once the skip values are calculated for the individual filters, step 908  
14 ascertains whether any of the calculated skip values exceed or equal the skip value  
15 tolerance. If none of the calculated skip values exceed or equal the skip value  
16 tolerance, then the method can return to step 902. Alternately, if the extrapolated  
17 lines are accurately predictable into the future, then the method can ascertain  
18 when, in fact, the calculated skip values will exceed or equal the skip value  
19 tolerance. If this is the case, or if the calculated skip values exceed or equal the  
20 skip value tolerance, then step 910 can provide instructions to synchronize the data  
21 streams. In the case where the calculated skip values actually exceed or equal the  
22 skip value tolerance, then the instructions to synchronize the data streams can be  
23 based on the actually calculated skip values. In the case where the method  
24 determines at which point in the future the calculated skip values will exceed or  
25

1 equal the skip value tolerance, then instructions can be based skip values that are  
2 calculated for the future.

### 3 4 Synchronizing Based on a Point of Reference

5 In another embodiment, a point of reference is defined and the data streams  
6 are periodically synchronized to the point of reference. Synchronization can take  
7 place periodically or when skip values exceed a defined skip value tolerance  
8 relative to the point of reference. As an example, consider Fig. 10.

9 There, a mapping of input timestamps to output timestamps for two data  
10 streams is shown generally at 1000. Assume that in this case, line 1002 represents  
11 that mapping for an audio stream and line 1004 represents the mapping for a video  
12 stream. Assume also that the requested playback rate is 2x. As shown, the audio  
13 stream is being provided with output timestamps such that it will be rendered  
14 slightly slower than the requested 2x rate. Likewise, the video stream is being  
15 provided with output timestamps such that it will be rendered slightly faster than  
16 the requested 2x rate. Over time, this disparity will lead to drifting between the  
17 streams which, in turn, will degrade the user experience.

18 Notice also in this example that a point of reference or "Reference Clock"  
19 is provided. This point of reference defines the reference to which the data  
20 streams are to be synchronized.

21 As in the above example, synchronization takes place by querying the  
22 filters for their input/output timestamp mappings and, if needed, their playback  
23 rate. The lines characterizing these mappings are then extrapolated to a selected  
24 input timestamp. In this example, assume that the filters associated with lines  
25 1002 and 1004 are queried when their input timestamp values are around 10 and



1 respond with output timestamps of 105 for the audio decoder and 95 for the video  
2 decoder. Assume that the ideal reference clock is at 100. Based on the  
3 information returned by the query, each of these lines is extrapolated to a selected  
4 input timestamp of 16 (as indicated by the dashed line). Now, skip values can be  
5 calculated based on the extrapolated lines and the point of reference. At 106, the  
6 audio timestamp would be 108.16 and the video timestamp would be 97.86. The  
7 reference clock would have advanced to 103. Here, the skip value for line 1002 is  
8 ascertained by, for an input timestamp value of 16, taking the difference of the  
9 output timestamps between the reference clock and line 1002 (i.e.  $103 - 108.16 = -$   
10  $5.16$ ). Similarly, the skip value for line 1004 is ascertained by, for an input  
11 timestamp value of 16, taking the difference of the output timestamps between the  
12 reference clock and line 1004 (i.e.  $103 - 97.86 = 5.14$ ). Now, the filters can be  
13 instructed to synchronize their data streams to the point of reference based on the  
14 calculated skip values. In the case of the filter processing the audio stream (i.e.  
15 corresponding to line 1002), the filter would subtract 5.16 from its output  
16 timestamp when its corresponding input time stamp value equals 16. Similarly, in  
17 the case of the filter processing the video stream (i.e. corresponding to line 1004),  
18 the filter would add 5.14 to its output timestamp when its corresponding input  
19 timestamp value equals 16.

20 In this way, the data streams can be synchronized to a point of reference or  
21 a reference clock. Synchronization can take place periodically or relative to a  
22 tolerance value that can be defined, as explained above.

1                    **Conclusion**

2                    The described methods and systems provide a general solution that can be  
3 applied to many multimedia streaming and network scheduling applications that  
4 utilize timestamps to render data streams. By synchronizing the data streams as  
5 described above, the user experience can be greatly enhanced. In addition,  
6 synchronization problems due to differing computation models as between  
7 different components that process data streams can be largely mitigated. This, in  
8 turn, can provide flexibility insofar as providing the ability to mix and match  
9 components that might, for example, be provided by different vendors.

10                  Although the invention has been described in language specific to structural  
11 features and/or methodological steps, it is to be understood that the invention  
12 defined in the appended claims is not necessarily limited to the specific features or  
13 steps described. Rather, the specific features and steps are disclosed as preferred  
14 forms of implementing the claimed invention.